Technical Report, March 2014
University of Ljubljana, Faculty of Computer and Information Science

arXiv:1403.7308v1 [stat.ML] 28 Mar 2014

Data generator based on RBF network

Marko Robnik-Šikonja
University of Ljubljana,
Faculty of Computer and Information Science,
Tržaška 25, 1001 Ljubljana, Slovenia
Marko.Robnik@fri.uni-lj.si

Abstract

There are plenty of problems where the data available is scarce and expensive. We propose a generator of semi-artificial data with similar properties to the original data which enables development and testing of different data mining algorithms and optimization of their parameters. The generated data allow a large scale experimentation and simulations without danger of overfitting. The proposed generator is based on RBF networks which learn sets of Gaussian kernels. Learned Gaussian kernels can be used in a generative mode to generate the data from the same distributions. To asses quality of the generated data we developed several workflows and used them to evaluate the statistical properties of the generated data, structural similarity and predictive similarity using supervised and unsupervised learning techniques. To determine usability of the proposed generator we conducted a large scale evaluation using 51 UCI data sets. The results show a considerable similarity between the original and generated data and indicate that the method can be useful in several development and simulation scenarios.

1 Introduction

One of technological challenges data analytics is facing is an enormous amount of data. This challenge is well known and recently a term ”big data” was coined with the purpose to bring attention to it and to develop new solutions. However, in many important application areas the excess of data is not a problem, quite the opposite, there just isn't enough data available. There are several reasons for this, the data may be inherently scarce (rare diseases, faults in complex systems, rare grammatical structures...), difficult to obtain (due to proprietary systems, confidentiality of business contracts, privacy of records...), expensive (obtainable with expensive equipment, requiring significant investment of human or material resources...), or the distribution of the events of interests is highly imbalanced (fraud detection, outlier detection, distributions with long tails...). For machine learning approaches the lack of data causes problems in

model selection, reliable performance estimation, development of specialized algorithms, and tuning of learning model parameters. While certain problems caused by scarce data are inherent to underrepresentation of the problem and cannot be solved, some aspects can be alleviated by generating artificial data similar to the original one. For example, similar artificial data sets can be of great help in tuning the parameters, development of specialized solutions, simulations, and imbalanced problems as they prevent overfitting of the original data set, yet allow sound comparison of different approaches.

Generating new data similar to a general data set is not an easy task. If there is no background knowledge available on the problem, we have to use the precious scarce data we posses to extract some of its properties and generate new semi-artificial data with similar properties. Weather this is acceptable in the context of the problem is not a matter of proposed approach, we assume that we can afford to set aside at least small part of the data for this purpose. This data may not be lost for modeling, but we shall be aware of extracted properties when considering possibility of overfitting.

The approaches used in existing data generators are limited to low dimensional data (up to 6 variables) or assume certain probability distribution, mostly normal; we review them in Sect. 2. Our approach is limited to classification problems. We first construct of a RBF network prediction model. RBF networks consist of Gaussian kernels which estimate probability density from training instances. Due to properties of Gaussian kernels (discussed in Section 3), the learned kernels can be used in a generative mode to produce new data. In such a way we overcome limitation to low dimensional spaces. We show that our approach can be successfully used for data sets with several hundred attributes and also with mixed data (numerical and categorical).

The paper is organized as follows. In Section 2 we review existing work on generating semi-artificial data. In Section 3 we present RBF neural networks and properties which allow us to generate data based on them. In Section 4 we present the actual implementation based on RSNNS package and explain details on handling nominal and numeric data. In Section 5 we discuss evaluation of generated data and its similarity to original data. We propose evaluation based on statistical properties of the data, as well as similarity between original and generated data estimated with supervised and unsupervised learning methods. In Section 6 we present the quality of the generated data and try to determine working conditions for proposed method as well as a suitable set of parameters. We shortly present an application of the generator for benchmarking of cloud bases big data analytics tool. In Section 7 we conclude with a summary, critical analysis and ideas for further work.

2 Related work

The area of data generators is full of interesting approaches. We cover only general approaches to data generation and do not cover methods specific for a certain problem or a class of problems.

The largest group of data generators is based on assumption about probability

distribution the generated data shall be drawn from. Most scientific computational engines and tools contain the random number generators for univariate data drawn from standard distributions. For example, R system [18] supports uniform, normal, log-normal, Student's t, F, Chi-squared, Poisson, exponential, beta, binomial, Cauchy, gamma, geometric, hypergeometric, multinomial, negative binomial, and Weibull distribution. Additional less-known univariate distribution-based random number generators are accessible through add-on packages. If we need univariate data from these distributions, we fit the parameters of the distributions and then use the obtained parameters to generate new data. For example, R package MASS [26] provides function fitdistr to obtain the parameters of several univariate distributions.

Random vector generators based on multivariate probability distributions are far less common. Effective random number generators exist for multivariate t and normal distribution with up to 6 variables. Simulating data from multivariate normal distribution is possible via a matrix decomposition of given symmetric positive definite matrix Σ containing variable covariances. Using the decomposed matrix and sequence of univariate normally distributed random variables one can generate data from multivariate normal distribution as discussed in Sect. 4. The approach proposed in this paper relies on the multivariate normal distribution data generator but does not assume that the whole data set is normally distributed. Instead it finds subspaces which can be successfully approximated with Gaussian kernels and use extracted distribution parameters to generate new data in proportion with the requirements.

To generate data from nonnormal multivariate distribution several transformational approaches have been proposed which start by generating data from a multivariate normal distribution and than transform it to the desired final distribution. For example, [24] proposes an iterative approximation scheme. In each iteration the approach generates a multivariate normal data that is subsequently replaced with the nonnormal data sampled from the specified target population. After each iteration, discrepancies between the generated and desired correlation matrices are used to update the intermediate correlation matrix. A similar approach for ordinal data is proposed by [5]. The transformational approaches are limited to low dimensional spaces where covariance matrix capturing data dependencies can be successfully estimated. In contrast, our method is not limited to specific data type. The problem space is split into subspaces where dependencies are more clearly expressed and subsequently captured.

Kernel density estimation is a method to estimate the probability density function of a random variable with a kernel function. The inferences about the population are made based on a finite data sample. Several approaches for kernel basedparameter estimation exist. The most frequently used kernels are Gaussian kernels. These methods are intended for low dimensional spaces with up to 6 variables [9].

An interesting approach to data simulation are copulas [17]. A copula is a multivariate probability distribution for which the marginal probability distribution of each variable is uniform. Copulas are estimated from the empirical observations and describe the dependence between random variables. They are based on Sklar's theorem that states that any multivariate joint distribution can be written with univariate

marginal distribution functions and a copula which describes the dependence structure between the variables. To generate new data one has to first select the correct copula family, estimate the parameters of the copula, and than generate the data. The process is not trivial and requires in-depth knowledge of the data being modeled. In principle the number of variables used in a copula is not limited, but in practice a careful selection of appropriate attributes and copula family is required [2, 15]. Copulas for both numeric and categorical data exist, but not for mixed types, whereas our approach is not limited in this sense.

3 RBF networks

RBF (Radial Basis Functions) networks have been proposed as a function approximation tool using locally tuned processing units, mostly Gaussian kernels [16, 30], but their development still continues [8, 29]. The network consists of three layers, see Figure 1 for an example. The input layer has a input units, corresponding to input features. The hidden layer contains kernel functions. The output layer consist of a single unit in case of regression or as many units as there are output classes in case of classifications. We assume a classification problem described with n pairs of $a-$dimensional training instances $(\mathbf{x}_i, y_i)$, where $\mathbf{x}_i \in \Re^a$ and y_i is one of class labels $1,2...C$. Hidden units computations in RBF network estimate the probability of each class y_u:

$$p(y_u|\mathbf{x}) = \sum_{j=1}^{k} c_j h_j(\mathbf{x}).$$

The weights c_j are multiplied by radial basis functions h_j, which are usually Gaussian kernels:

$$h_j(\mathbf{x}) = \exp(-\frac{||\mathbf{x}-\mathbf{t}_j||}{\sigma_j^2}).$$

Vectors $\mathbf{t}_j$ present k centers and σ_j are widths of the kernels. The centers and kernel widths σ_j have to be learned or set in advance. The kernel function h is applied to the Euclidian distance between each center $\mathbf{t}_j$ and given instance $\mathbf{x}$. Kernel functions have a maximum at zero distance from the center, while the activation is close to zero for instances which are further away from the center.

Most algorithms used to train RBF networks require a fixed architecture in which the number of units in the hidden layer must be determined before the training starts. To avoid manual setting of this parameter and to automatically learn kernel centers $\mathbf{t}_j$, weights c_j, and standard deviations σ_j, several solutions have been proposed [20, 4], among them RBF with Dynamic Decay Adjustment (DDA)[4] which we use in this work. The RBF DDA builds a network by incrementally adding an appropriate number of RBF units. Each unit encodes instances of only one class. During the process of adding new units the kernel widths σ_j are dynamically adjusted (decayed) based on information about neighbors. RBFs trained with the DDA algorithm often achieve

classification accuracy comparable to Multi Layer Perceptrons (MLPs) but training is significantly faster [4, 30].

An example of RBF-DDA network for classification problem with 4 features and a binary class is presented in Fig. 1. The hidden layer of RBF-DDA network contains Gaussian units, which are added to this layer during training. The input layer is fully connected to the hidden layer. The output layer consists of one unit for each possible class. Each hidden unit encodes instances of one class and is therefore connected to exactly one output unit. For classification of a new instance a winner-takes-all approach is used, i.e. the output unit with the highest activation determines the class value.

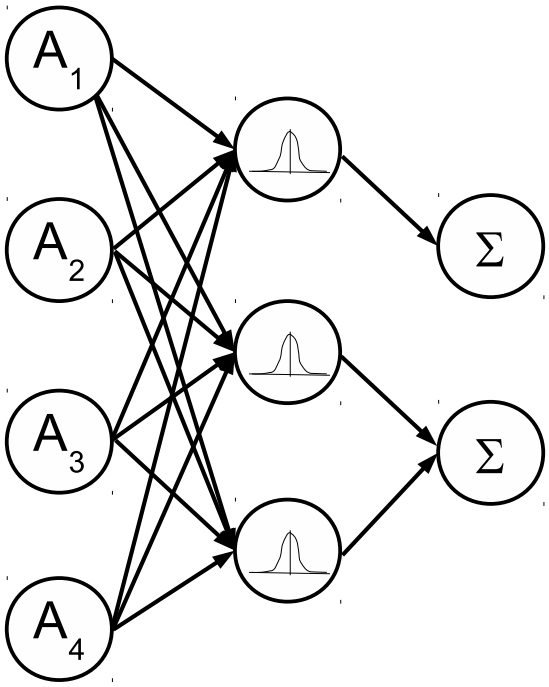

Figure 1: Structure of RBF-DDA network for classification problem with 4 attributes, 3 hidden units, and a binary class.

Our data generator uses the function rbfDDA implemented in R package RSNNS [3] which is a R port of SNNS software [30]. The implementation uses two parameters: a positive threshold Θ^+ and a negative threshold Θ^- as illustrated on Fig. 2. The two thresholds define an upper and lower bound for the activation of training instances. Default values of thresholds are $\Theta^+ = 0.4$ and $\Theta^- = 0.2$. The thresholds define a safety area where no other center of a conflicting class is allowed. In this way a good separability of classes is achieved. In addition, each training instance has to be in the inner circle of at least one center of the correct class.

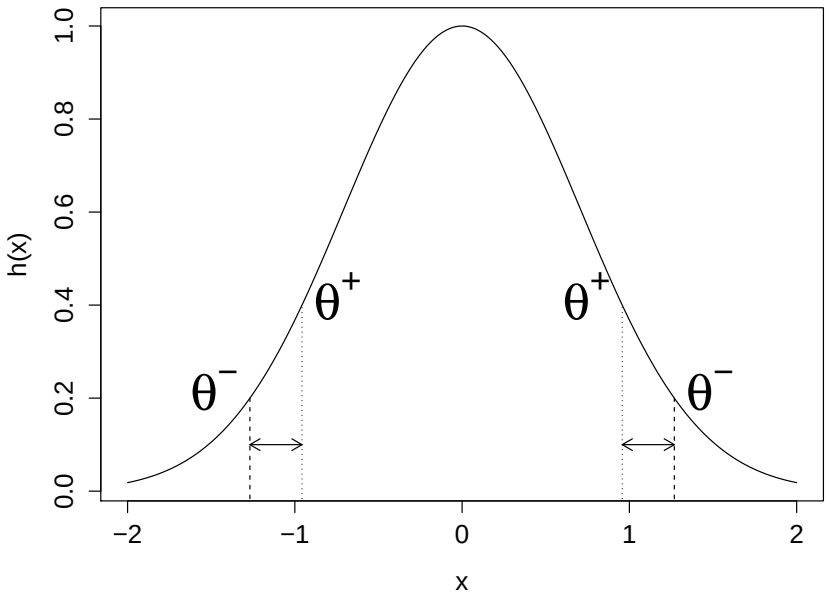

Figure 2: The thresholds Θ^+ and Θ^- illustrated on a single Gaussian unit. The arrows indicate the conflicting areas where no center of a different class can be placed.

4 Data generator

The idea of the proposed data generation scheme is to extract local Gaussian kernels from the learned RBF-DDA network and generate data from each of them in proportion to the desired class value distribution. When class distribution different from the empirically observed is desired, the distribution has to be specified as an input parameter.

A notable property of a Gaussian kernels is their ability not to be used only as discriminative models but also as generative models. To generate data from multivariate normal distribution $N(\mu,\Sigma)$ one can exploit the following property of multivariate Gaussian distribution:

$$\text{if } \mathbf{X} \sim N(\mu,\Sigma), \text{ then } \mathbf{Y} = A\mathbf{X}+\mathbf{b} \sim N(A\mu+\mathbf{b},\ A\Sigma A^T). \tag{1}$$

When we want to simulate multidimensional $\mathbf{Y} \sim N(\mu,\Sigma)$, for a given symmetric positive definite matrix Σ, we first construct a sample $\mathbf{X} \sim N(\mathbf{0},1)$ of the same dimensionality. The $\mathbf{X} \sim N(\mathbf{0},1)$ can easily be constructed using independent variables $X_i \sim N(0,1)$. Next we decompose $\Sigma = AA^T$ (using Choleski or eigenvalue decomposition). With the obtained matrix A and $\mathbf{X}$ we use Eq. (1) to get

$$\mathbf{Y} = A\mathbf{X}+\mu \sim N(A\mathbf{0}+\mu, A1A^T) = N(\mu,\Sigma).$$

In our implementation we use function mvrnorm from R package MASS [26] which decomposes covariance matrix Σ with eigenvalue decomposition due to better stability [21].

4.1 Construction of generator

The pseudo code of the proposed generator is given in Figure 3. The input to the generator is the available data set and two parameters. The parameter *minW* controls the minimal acceptable kernel weight. The weight of the kernel is defined as the number of training instances which achieve maximal activation with that kernel. All the learned kernels with weight less than *minW* are discarded by data generator to prevent overfitting of the training data. The boolean parameter *nominalAsBinary* controls the treatment of nominal attributes as described in Sect. 4.2.

Due to specific demands of RBF-DDA algorithm the data has to be preprocessed first (line 2 in Fig. 3). This preprocessing includes normalization of attributes to $[0,1]$ and preparation of nominal attributes (see Sect. 4.2). Function rbfPrepareData returns normalized data D and normalization parameters N, which are used later when generating new instances. The learning algorithm takes the preprocessed data and returns the classification model M in the form of Gaussian kernels (line 3). We store the learned parameters of the Gaussian kernels, namely their centers $\mathbf{t}$, weights w, and class values c (lines 4, 5, and 6). The kernel weight w_k equals the proportion of training instances which are activated by the k-th Gaussian unit. The class value c_k of the unit corresponds to the output unit connected to the Gaussian unit k (see Fig. 1 for an illustration). Theoretically, this extracted information would be sufficient to generate new data, however there are several practical considerations, which have to be taken into account if one is to generate new data comparable to the original one.

Input: data set $D = \{(\mathbf{x}_i, y_i)_{i=1}^n\}$, parameters *minW*, *nominalAsBinary*
Output: a list L of Gaussian kernels and a list N of attribute normalization parameters

1 **Function** rbfGen(*D*, *minW*, *nominalAsBinary*)
  // preprocess the data to get $[0,1]$ normalized data D and normalization
2 $(D,N) \leftarrow$ rbfPrepareData(*D*, *nominalAsBinary*) // parameters N
3 $M \leftarrow$ rbfDDA(D) // learn RBF model consisting of kernels
4 **foreach** kernel $k \in M$ **do**
5 **if** $w_k \geq minW$ **then** // store only kernels with sufficient weight
6 $L_k \leftarrow (\mathbf{t}_k, w_k, c_k)$ // store center, weight, and class
7 **for** $i \in 1 \ldots n$ **do** // find activation unit of each instance
8 $z_i \leftarrow \arg\max_{k \in M} \exp(\frac{-||\mathbf{x_i} - \mathbf{t_k}||}{\sigma_k^2})$
9 **foreach** kernel $k \in M$ **do** // estimate empirical kernel width
10 $\Sigma_k \leftarrow \text{std}(\{\mathbf{x}_i; z_i = k\})$ // compute spread on matching instances
11 $L_k \leftarrow L_k \cup \Sigma_k$ // add Σ_k to list item L_k
12 **return** *(L, N)*

Figure 3: The pseudo code of creating a RBF based data generator.

The task of RBF-DDA is to discriminate between instances with different class values, therefore widths of the kernel are set during the learning phase in such a way that

majority of instances are activated by exactly one kernel. Widths of the learned kernels therefore prevent overlapping of competing classes. For the purpose of generating new data the with of the kernel shall be different (not so narrow), or we would only generate instances in the near proximity of kernel centers i.e. existing training instances. The approach we adopted is to take the training instances that activate the particular kernel (lines 7 and 8) and estimate their empirical variance (lines 9, 10, and 11) in each dimension, which is later, in the generation phase, used as the width of the Gaussian kernel. The Σ matrix extracted from the network is diagonal, with elements presenting the spread of training instances in each dimension. The algorithm returns the data generator consisting of the list of kernel parameters L and normalization parameters N (line 12).

4.2 Preprocessing the data

Function rbfPrepareData does three tasks: it imputes missing values, prepares nominal attributes, and normalizes the data. The pseudo code of data preprocessing is in Fig. 4.

Input: data set $D = \{(\mathbf{x}_i, y_i)_{i=1}^n\}$, parameter *nominalAsBinary*
Output: preprocessed data D', a list T with information on attribute transformations

```
1  Function rbfPrepareData(D, nominalAsBinary)
2  for j ∈ 1...a do // preprocessing of attributes
3  |  D'(x_.j) ← imputeMissing(x_.j) // imputation of missing values
4  |  if (isNominal(x_.j)) // encode nominal attributes
5  |  |  if (nominalAsBinary)
6  |  |  |  D' ← encodeBinary(x_.j)
7  |  |  else
8  |  |  |  D' ← encodeInteger(x_.j)
9  |  D'(x_.j) ← (x_.j − min x_.j) / (max x_.j − min x_.j)  // normalize attributes to [0,1]
   |  // store normalization and encoding parameters
10 |  T_j ← (min x_.j, max x_.j − min x_.j, encoding(x_.j) )
11 D' ← class y with binary encodings
12 return (D', T)
```

Figure 4: Preprocessing the data for RBF-DDA algorithm; $x_{.j}$ stands for values of attribute j.

The rbfDDA function in R does not accept missing values, so we have to impute them (line 3). While several advanced imputation strategies exist, the classification accuracy is not of the uttermost importance in our case, so we resorted to median based imputation for numeric attributes, while for nominal attributes we use the most frequent category.

Gaussian kernels are defined only for numeric attributes, so rbfDDA treats all the attributes, including nominal, as numeric. Each nominal attribute is converted to numeric (lines 4-8). We can simply assigning each category a unique integer from 1 to the number of categories (line 8). This may be problematic as this transformation has established an order of categorical values in the converted attribute, inexistent in the original attribute. For example, for attribute *Color* the categories $\{red, green, blue\}$ are converted into values $\{1,2,3\}$, respectively, meaning that the category *red* is now closer to *green* than to *blue*. To solve this problem we use the binary parameter *nominalAsBinary* (line 5) and encode nominal attributes with several binary attributes when this parameter is set to *true* (line 6). Nominal attributes with more than two categories are encoded with the number of binary attributes equal to the number of categories. Each category is encoded by one binary attribute. If the value of the nominal attribute equals the given category, the value of the corresponding binary attribute is set to 1, while the values of the other encoding binary attributes equal 0. E.g., *Color* attribute with three categories would be encoded with three binary attributes $C_{red}, C_{green}, C_{blue}$. If the value of the attribute is $Color = green$ then the binary encoding of this value is $C_{red} = 0, C_{green} = 1, and C_{blue} = 0$. The same binary encoding is required also for class values (line 11).

The rbfDDA function in R expects data to be normalized to $[0,1]$ (line 9). As we want to generate new data in the original, unnormalized form, we have to store the computed normalization parameters (line 10) and, together with attribute encoding information, pass them back to the calling rbfGen function.

4.3 Generating new data

Once we have a generator (produced by function rbfGen) , we can use it to generate new instances. By default the method generates with class values proportionally to the number of class values in the training set of the generator, but the user can specify the desired class distribution as a parameter p.

A data generator consists of a list L of parameters describing Gaussian kernels and information on attribute transformations T. Recall that information for each kernel k contains the location of kernel's center $\mathbf{t}_k$, weight of kernel w_k, class value c_k, and estimated standard deviation Σ_k. An input to newdata function are also parameters *size* specifying the number of instances to be generated, p the desired distribution of class values, *var* controlling the width of the kernels, and $defaultSpread$ as the width of the kernel if estimated width is 0.

Function starts by creating an empty data set D (line 2) and than generates instances with each of the kernels stored in the kernel list L (lines 2-11).The weight of the kernel k, the desired class probability p_k, and the overall number of instances to be generated *size* determine the number of instances g to be generated with each kernel (line 4). The weight of the kernel is normalized with the weights of the same class kernels $\frac{w_k}{\sum_{i=1}^{|L|} w_i \cdot I(c_i = c_k)}$, where $I()$ presents an indicator function. The width of the kernel determines the spread of the generated values around the center. By default we use the

Input: L - a list of Gaussian kernels, T - an information on attribute normalization and encoding, *size* - the number of instances to be generated, p - a vector of desired class distribution, *var* - a parameter controlling the width of kernels, $defaultSpread$ - the width of the kernel if estimated width is 0

Output: new data set $F = \{(\mathbf{x}_i, y_i)_{i=1}^{size}\}$

```
1  Function newdata(L, T, size, p, var, defaultSpread)
2  D ← {} // create an empty temporary data set
3  foreach kernel k ∈ L do
4  |  g ← (w_k / Σ_{i=1}^{|L|} w_i·I(c_i=c_k)) · p_{c_k} · size // number of instances to generate
   |  // set kernel width
5  |  if (var="estimated") then Σ = Σ_k with zeros substituted by defaultSpread;
6  |  else if (var="Silverman") then Σ = silverman(Σ_k, n, a) ; // heuristic rule
7  |  H ← mvrnorm(n=g, mu=t_k, Sigma=Σ) // generate new data with kernel k
8  |  H ← makeConsistent(H, T_k) // check and fix inconsistencies
9  |  H(y) ← c_k // assign class value from the kernel
10 |  D ← D ∪ H // append generated data to D
11 for j ∈ 1...a do // transform attributes back to original scales and encodings
12 |  if (T_j.nominal) then // decode nominal attributes
13 |  |  if (T_j.binaryEncoded) then
14 |  |  |  F(x_.j) ← decodeBinary(D(x_.j), T_j)
15 |  |  else
16 |  |  |  F(x_.j) ← decodeInteger(D(x_.j), T_j)
17 |  else
18 |  |  F(x_.j) ← D(x_.j) · T_j.span + T_j.min // denormalize attributes
19 return F
```

Figure 5: The pseudo code for creating new instances with RBF generator.

spread as estimated from the training data (line 5). Zeros in individual dimensions are optionally replaced by value of parameter $defaultSpread$. For kernel width it is also possible to use the generalization of Silverman's rule of thumb for multivariate case (line 6) [9]. In this case the covariance matrix used is diagonal, i.e., diag$(w_1, \ldots w_a)$, and kernel width in each dimension is set to

$$w_j = \left(\frac{4}{a+2}\right)^{\frac{1}{(a+4)}} n^{\frac{-1}{a+4}} \sigma'_j,$$

where n is the sample size (in our case number of training instances that activate the particular kernel, and σ'_j is the estimated spread in that dimension.

The data is generated by mvrnorm function (line 7). The function takes as input g, the number of instances to generate, the center of the kernel $\mathbf{t}_k$, and the diagonal

covariance matrix Σ_k. Function exploits the property of Gaussian kernels from Eq. (1) and decomposes covariance matrix Σ with eigenvalue decomposition. The generated data has to be checked for consistency (line 8), i.e., generated attribute values have to be in $[0,1]$ interval, nominal attributes have to be rounded to values encoding existing categories, etc. As some instances are rejected during this process in practice we generate more than g instances with mvrnorm but retain only the desired number of them. We assign the class value to the generated instances (line 9) and append them to D (line 10).

When the data are generated with all kernels we have to transform the generated instances to the original scale and encodings. For each nominal attribute we check its encoding (either as a set of binary attributes or as an integer), and transform it back to the original form (lines 11-18). Numeric attributes are denormalized and transformed back to the original scale using minimums and spans stored in T (line 18). The function returns the generated data set (line 19).

4.4 Visual inspection of generated data

As a demonstration of the generator we graphically present the generated data on two simple data sets. The first data set forms a two dimensional grid where attributes A_1 and A_2 are generated with three Gaussian kernels with centers at $(-5,-5)$, $(0,0)$, and $(5,5)$. Each group of 500 instances is assigned a unique class value (red, blue, and green, respectively) as illustrated in Fig 6a. The generator based on this data consists of eight Gaussian kernels (two for red and blue class each, and four for green class). We illustrate 1500 instances generated with this generator in Fig 6b. As the rbfDDA learner did not find the exact locations of the original centers it approximated the data with several kernels, so there is some difference between the original and generated data, but individual statistics are close as shown in Table 1.

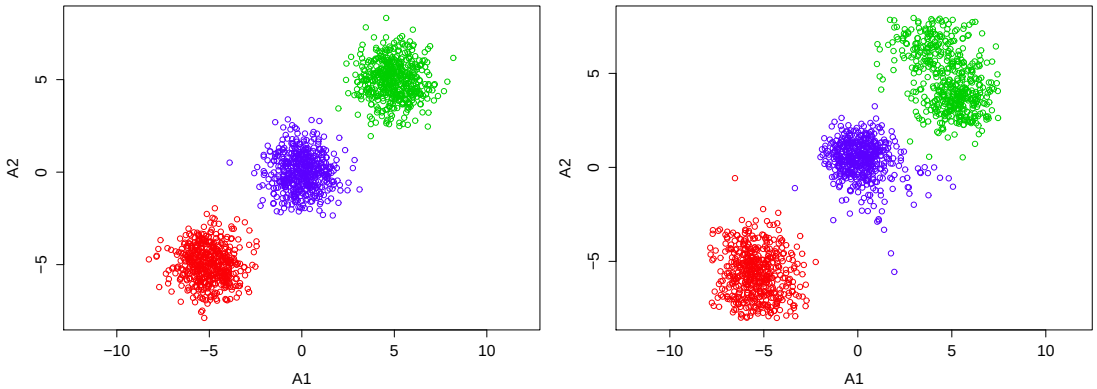

Figure 6: An illustration of generated data on a simple two dimensional dataset with the original data on the left-hand side and generated data on the right-hand side.

Another simple example is the well known Iris data set which consists of 50 samples from each of three species of Iris (Iris setosa, Iris virginica and Iris versicolor).

Table 1: The summary of the original and generated data sets from Fig. 6.

	original data		generated data	
value	A_1	A_2	A_1	A_2
Minimum	-8.27	-7.88	-7.85	-8.00
1st Quartile	-4.34	-4.28	-4.74	-4.83
Median	0.00	0.00	0.13	0.54
Mean	-0.02	0.02	-0.14	-0.18
3rd Quartile	4.32	4.26	4.00	3.42
Maximum	8.19	8.33	7.46	7.95

Four features were measured from each sample: the length and the width of the sepals and petals, in centimeters. The scatter plots of the original data sets are shown in the Fig. 7a where class values are marked with different colors. The generator based on this data consisting of 31 Gaussian units generated 150 instances shown in Fig. 7b. The graphs show considerable similarity between matching pairs of scatter plots.

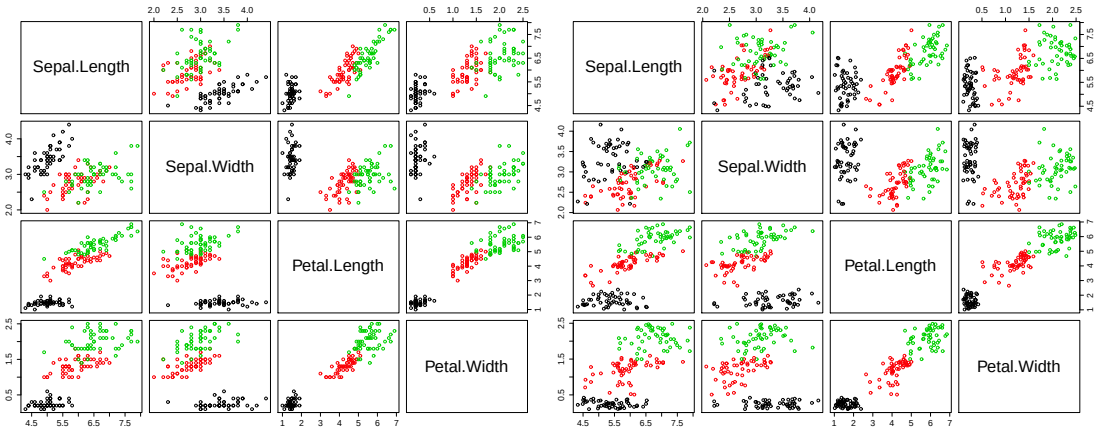

a) b)

Figure 7: A visual comparison of original (left-hand side) and generated data (right-hand side) for the well known Iris data set.

5 Data quality

We are not aware of any other data generator capable of generating data similar to existing data sets with no limitations in the number and type of attributes. The quality of existing data generators is mostly evaluated by comparing standard statistics: mean, median, standard deviation, skewness, and kurtosis. While these statistics are important indicators of the quality of generated data, they are insufficient for data sets with more attributes (e.g., more than 6). They are computed for each attribute separately, thereby not presenting an overall view of the data, do not take possible interactions

between attributes into account, and are difficult to compare when the number of attributes increases. These statistics may also not convey any information about how appropriate and similar is the generated data for machine learning and data mining tasks. To resolve this difficulties and quantify similarity between original and generated data we developed several data quality measures described below. We use measures incorporating standard statistics, measures based on clustering and measures based on classification performance.

5.1 Standard statistics

Standard statistics for numeric attributes we use are the mean, standard deviation, skewness, and kurtosis. We compare also value distributions of attributes from original and generated data. For this we use Hellinger distance and Kolmogorov-Smirnov test (KS). The workflow of the whole process is illustrated in Fig 8.

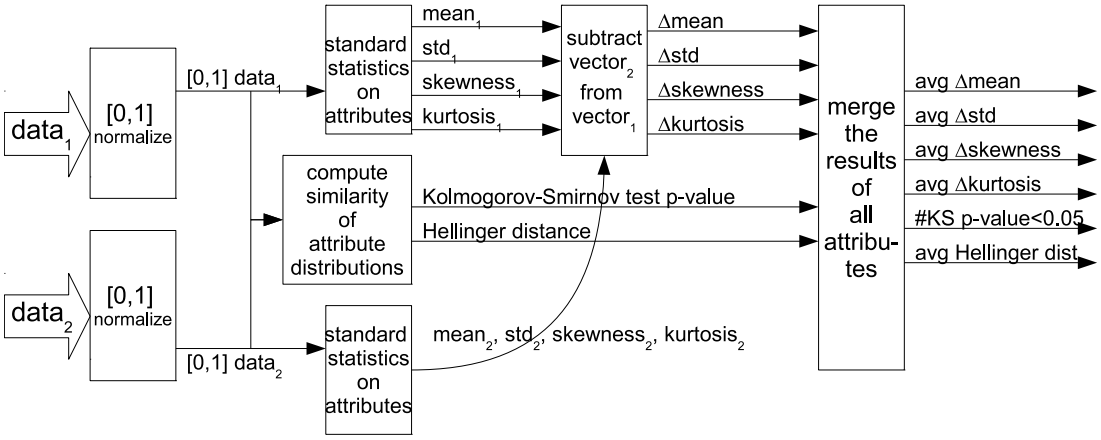

Figure 8: The workflow of comparing standard statistics between two data sets.

We normalize each numeric attribute to $[0,1]$ to make comparison between attributes sensible. Input to each comparison are two data sets (original and generated). Comparison of attributes' statistics computed on both data sets is tedious especially for data sets with large number of attributes. We therefore first compute standard statistics on attributes and then subtract statistics of the second data set from statistics of the first data set. To summarize the results we report only the average difference for each of the statistics. To compare distributions of attribute values we use Hellinger distance for discrete attributes and KS test for numerical attributes. The Hellinger distance between two discrete univariate distributions $P = (p_1, \ldots p_k)$ and $Q = (q_1, \ldots q_k)$ is defined as

$$H(P,Q) = \frac{1}{\sqrt{2}} \sqrt{\sum_{i=1}^{k} (\sqrt{p_i} - \sqrt{q_i})^2}$$

The maximal Hellinger distance between two distributions is 1. For numerical attributes we use two sample KS test, which tests whether two one-dimensional proba-

bility distributions differ. The statistics uses maximal difference between two empirical cumulative distribution functions $F_1(x,n_1)$ and $F_2(x,n_2)$, on samples of size n_1 and n_2, respectively:

$$D(n_1,n_2) = \sup_x |F_1(x,n_1) - F_2(x,n_2)|.$$

To get an overall picture we again report only the average Hellinger distance over all discrete attributes, and percentage of numeric attributes for which p-value of KS test was below 0.05. We set the null hypothesis that that attributes' values in both data sets are drawn form the same distribution. While these averages do not have a strict statistical meaning, they do illustrate the similarity of two data sets.

5.2 Comparing clustering

Cluster information is an important introspection into the structure of the data. We compare similarity of the clusterings obtained for two data sets (original and generated). To estimate the similarity based on clusterings we use the Adjusted Rand Index (ARI)[11].

5.2.1 Similarity of two clusterings

Starting from the data set $D = \{(\mathbf{x}_i)_{i=1}^n\}$ with n data points, we assume two different clusterings of D, namely $U = \{U_1, U_2, ..., U_u\}$ and $V = \{V_1, V_2, ..., V_v\}$, where $U_1 \cap U_2 \cap \cdots \cap U_u = \emptyset$, $U_1 \cup U_2 \cup \cdots \cup U_u = D$, $V_1 \cap V_2 \cap \cdots \cap V_v = \emptyset$, $V_1 \cup V_2 \cup \cdots \cup V_v = D$. The information on overlap between clusters of U and V can be expressed with $u \times v$ contingency table as shown in Table 2.

Table 2: The contingency table of clusterings overlap $n_{i,j} = |U_i \cap V_j|$.

U/V	V_1	V_2	$\ldots$	V_v	sum
U_1	$n_{1,1}$	$n_{1,2}$	$\ldots$	$n_{1,v}$	u_1
U_2	$n_{2,1}$	$n_{2,2}$	$\ldots$	$n_{2,v}$	u_2
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
U_u	$n_{u,1}$	$n_{u,2}$	$\ldots$	$n_{u,v}$	u_u
sum	v_1	v_2	$\ldots$	v_v	n

There are several measures comparing clusterings based on counting the pairs of points on which two clusterings agree or disagree [28]. Any pair of data points from the total of $\binom{n}{2}$ distinct pairs in D falls into one of the following 4 categories.

- N_{11}, the number of pairs that are in the same cluster in both U and V;
- N_{00}, the number of pairs that are in different clusters in both U and V;
- N_{01}, the number of pairs that are in the same cluster in U but in different clusters in V;

- N_{10}, the number of pairs that are in different clusters in U but in the same cluster in V.

The values N_{11}, N_{00}, N_{01}, and N_{10} can be computed from contingency table [11]. Values N_{11} and N_{00} indicate agreement between clusterings U and V, while values N_{01} and N_{10} indicate disagreement between U and V. The original Rand Index [19] is defined as

$$RI(U,V) = \frac{N_{00} + N11}{\binom{n}{2}}.$$

The Rand Index lies between 0 and 1 and takes the value 1 when the two clusterings are identical, and the value 0 when no pair of points appear either in the same cluster or different clusters in both U and V. It is desirable that a similarity indicator would take value close to zero for two random clusterings, which is not true for RI. The Adjusted Rand Index [11] fixes this by using generalized hypergeometric distribution as a model of randomness and computes expected number of entries in the contingency table. It is defined as

$$ARI = \frac{RI - E[RI]}{\max RI - E[RI]} = \frac{\sum_{i=1}^{u}\sum_{j=1}^{v}\binom{n_{i,j}}{2} - [\sum_{i=1}^{u}\binom{u_i}{2}\sum_{j=1}^{v}\binom{v_j}{2}]/\binom{n}{2}}{\frac{1}{2}[\sum_{i=1}^{u}\binom{u_i}{2} + \sum_{j=1}^{v}\binom{v_j}{2}] - [\sum_{i=1}^{u}\binom{u_i}{2}\sum_{j=1}^{v}\binom{v_j}{2}]/\binom{n}{2}}. \quad (2)$$

The ARI has expected value of 0 for random distribution of clusters, and value 1 for perfectly matching clusterings. ARI can also be negative.

5.2.2 A workflow for comparing clusterings on two data sets

The ARI is used to compare two different clusterings on the same set of instances, while we want to compare similarity of two different sets of instances. To overcome this obstacle, we cluster both data sets separately and extract medoids of the clusters for each clustering. The medoid of a cluster is an existing instance in the cluster whose average similarity to all instances in the cluster is maximal. For each instance in the first data set, we find the nearest medoid in the second clustering and assign it to that cluster, thereby getting a joint clustering of both data sets based on the cluster structure of the second data set. We repeat the analogous procedure for the second data set and get a joint clustering based on the first data set. These two joint clusterings are defined on the same set of instances (union of both original and generated data), therefore we can use ARI to asses similarity of the clusterings and compare structure of both data sets. The workflow of cluster based comparison of two data sets is illustrated in Fig. 9.

As we need to assign new instances to existing clustering we selected partitioning with medoids (PAM) clustering algorithm [12], which, besides partitions, outputs also medoids. Distance to the medoids is the criterion we use to assign new instances to existing clusters. PAM clustering is implemented in R package cluster [14]. To use this method we first computed distances between instances of each data set using Gower's method [7]. The method normalizes numeric attributes to $[0,1]$ and uses 0-1 scoring

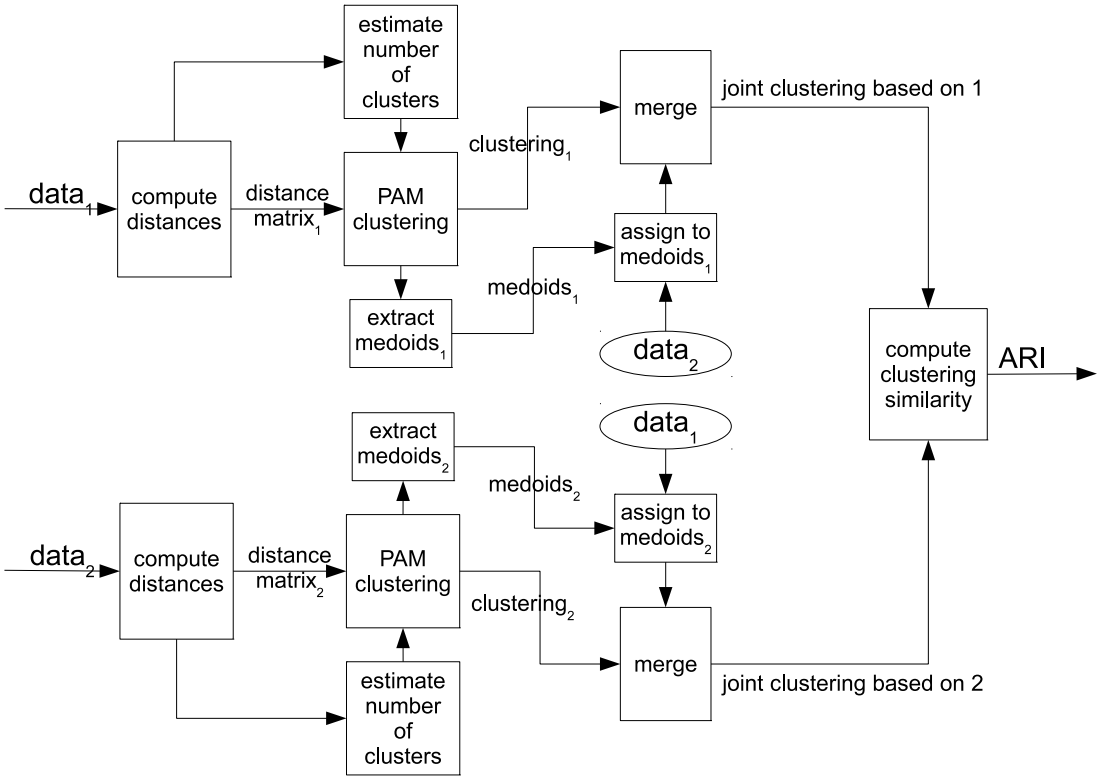

Figure 9: The workflow of comparing two data sets based on clustering similarity.

of dissimilarity between nominal attributes (0 for the same, 1 for different categories). The distance is a sum of dissimilarities over all attributes.

The number of clusters is set to an estimated optimal value separately for original and generated data set. The number of clusters is estimated with optimum average silhouette width method and is computed in R package fpc [10]. ARI is computed with R package mclust [6].

5.3 Comparing classification

The classification is probably the most important task in machine learning and data mining. The judgment how good substitute for original data set the generated instances are, is therefore largely dependent on classification similarity between data sets. The scenario we propose to measure similarity of classification performance is shown in Fig. 10.

The basic idea is to train models with the original data and with the generated data. Both models are tested on yet unseen original and generated data and the performances are compared. If the performance of a model trained on original data is comparable for original and generated data this is an indicator that the generated data is within the original distribution (i.e., there are no significant outliers and all the aspects of the original data are captured). If performance of a model trained on the generated data is

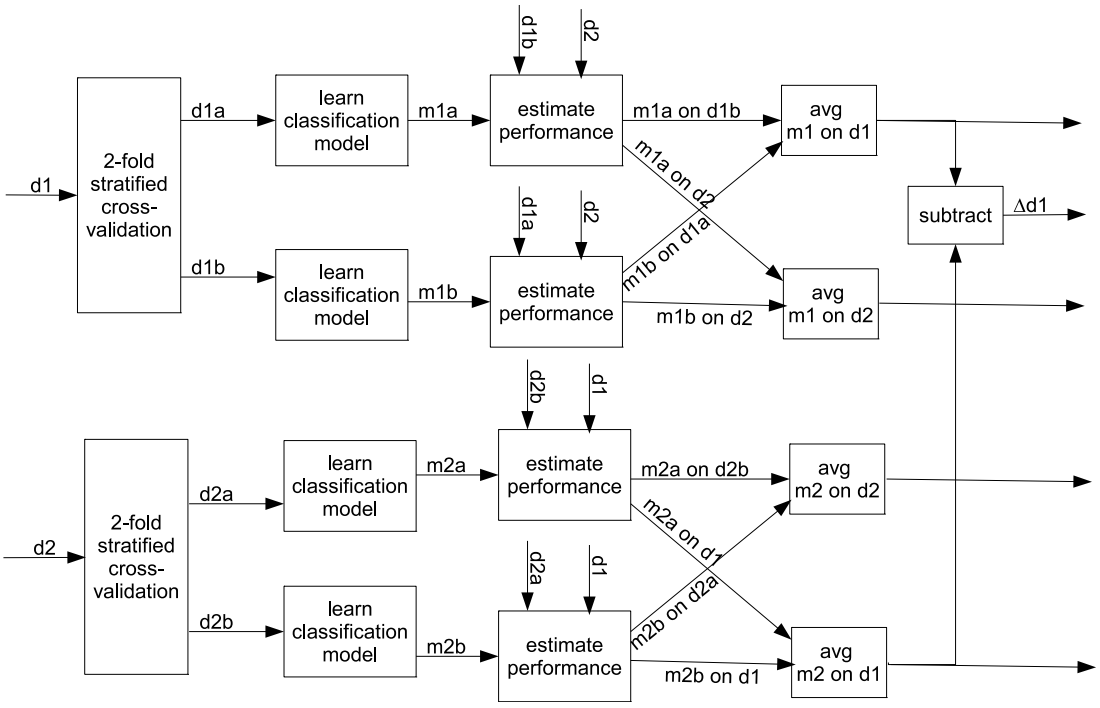

Figure 10: The workflow for comparing two data sets based on classification performance.

comparable for original and generated data this shows that the generated data enables comparable learning and has a good coverage of the original distribution, therefore the generator is able to produce good substitutes for original data concerning machine learning and data mining. Additionally, if the model trained on original data achieves better performance on the generated data than on original data, this indicates that the generator is oversimplified and does not cover all peculiarities of the original data.

In our testing workflow (Fig. 10) we start with two data sets d1 and d2 (e.g. original and generated one) and split them randomly but stratified into two halves (d1 produces d1a and d1b, d2 is split into d2a and d2b). Each of the four splits is used to train a classifier, and we name the resulting models m1a, m1b, m2a, and m2b, respectively. We evaluate the performance of these models on data unseen during training, so m1a is tested on d1b and d2, m1b is tested on d1a and d2, m2a uses d2b and d1, and m2b uses d2a and d1 as the testing set. Each test produces a performance score (e.g., classification accuracy, AUC...) which we can average as in a 2-fold cross-validation to get the following estimates:

- performance of m1 on d1 (model built on original data and tested on original data) is an average performance of m1a on d1b and m1b on d1a,
- performance of m1 on d2 (classifier built on original data and tested on generated data) is an average performance of m1a on d2 and m1b on d2,
- performance of m2 on d2 (model built on generated data and tested on generated

data) is an average performance of m2a on d2b and m2b on d2a,

- performance of m2 on d1 (model built on generated data and tested on original data) is an average performance of m2a on d1 and m2b on d1.

These estimates already convey an important information as discussed above, but we can subtract performances of m1 on d1 and m2 on d1 (models built on original and generated data, both tested on original data) to get Δd1. This difference, in our opinion, is the most important indicator how suitable is the generated data for development of classification methods. If, in ideal case, this difference would be close to zero, the generated data would be a good substitute for the lack of original data, as we can expect that performance of developed methods will be comparable when used on original data.

6 Evaluation

We want to verify if the proposed generator produces data which is consistent with the original and if it covers the whole space the original data set does. We try to determine working conditions of the generator: on which data set it works and where it fails, on what sort of problems it veritably reproduces the original and where it is less successful. We first describe the evaluation scenario and compare original and generated data. Afterwards we examine parameters of the generator and propose reasonable defaults.

To evaluate the generator we performed a large scale empirical evaluation using 51 data sets from UCI repository [1] with great variability in the number of attributes, types of attributes and number of class values. We used R package readMLData [25] which provides an uniform interface for manipulation of UCI data sets. Assuming that one would mostly desire to generate semi-artificial data when the number of original instances is rather small and to keep computational load of the evaluation low, we limited the number of original instances to be between 50 and 1000 (lower limit is necessary to assure sufficient data for both the generator and the testing set). Taking these conditions into account we extracted 51 classification data sets from a collection of 92 data sets kindly provided by author of the package readMLData. The characteristics of these data sets are presented in Table 3.

For each data set a generator based on rbfDDA learner was constructed with function rbfGen (Fig. 3). We used the value of parameter $minW = 1$ for all data sets and compared both variants of encoding for nominal attributes. The produced generator was used to generate the same number of instances as in the original data set with the same distribution of classes using the function newdata (Fig. 5). The width of the kernels was estimated from the training instances by setting the parameter var="estimated".

We compared the original data set with the generated one using the three workflows described in Sect. 5.

Statistics of attributes: we compared standard statistics of numeric attributes (Fig. 8) - mean. standard deviation, skewness, and kurtosis. The interpretations of

Table 3: The characteristics of data sets used. The column are: *n* - number of instances, *a* - number of attributes, numeric - number of numeric attributes, discrete - number of discrete attributes, v/a - average number of values per discrete attribute, *C* - number of class values, majority % - proportion of majority class in percent, missing % - percentage of missing values.

dataset	*n*	*a*	numeric	discrete	v/a	*C*	majority %	missing %
annealing	898	38	6	32	2.4	5	76.2	0.00
arrhythmia	452	279	206	73	1.9	13	54.2	0.32
audiology	226	69	0	69	2.2	24	25.2	2.03
automobile	205	25	15	10	6.0	6	32.7	1.15
balance-scale	625	4	4	0	0.0	3	46.1	0.00
breast-cancer	286	9	0	9	4.6	2	70.3	0.35
breast-cancer-wdbc	569	30	30	0	0.0	2	62.7	0.00
breast-cancer-wisconsin	699	9	9	0	0.0	2	65.5	0.25
bridges.version1	106	11	4	7	3.0	7	41.5	5.57
bridges.version2	106	11	1	10	3.2	7	41.5	5.57
bupa	345	6	6	0	0.0	2	58.0	0.00
credit-screening	690	15	6	9	4.4	2	55.5	0.64
cylinder-bands	540	37	20	17	8.7	2	57.8	5.00
dermatology	366	34	1	33	3.9	6	30.6	0.06
ecoli	336	7	7	0	0.0	8	42.6	0.00
flags	194	28	2	26	4.5	8	35.6	0.00
glass	214	9	9	0	0.0	6	35.5	0.00
haberman	306	3	2	1	12.0	2	73.5	0.00
heart-disease-cleveland	303	13	6	7	2.7	5	54.1	0.15
heart-disease-hungarian	294	13	6	7	2.7	2	63.9	20.46
hepatitis	155	19	6	13	2.0	2	79.4	5.67
horse-colic	368	21	7	14	3.7	2	63.0	24.90
house-votes-84	435	16	0	16	3.0	2	61.4	0.00
ionosphere	351	34	34	0	0.0	2	64.1	0.00
iris	150	4	4	0	0.0	3	33.3	0.00
labor-negotiations	57	16	8	8	2.6	2	64.9	35.74
lymphography	148	18	3	15	2.9	4	54.7	0.00
monks-1	556	6	0	6	2.8	2	50.0	0.00
monks-2	601	6	0	6	2.8	2	65.7	0.00
monks-3	554	6	0	6	2.8	2	52.0	0.00
pima-indians-diabetes	768	8	8	0	0.0	2	65.1	0.00
post-operative	90	8	1	7	2.7	3	71.1	0.41
primary-tumor	339	17	0	17	2.2	21	24.8	3.90
promoters	106	57	0	57	4.0	2	50.0	0.00
sonar.all	208	60	60	0	0.0	2	53.4	0.00
soybean-large	683	35	0	35	2.8	19	13.5	9.77
spect-SPECT	267	22	0	22	2.0	2	79.4	0.00
spect-SPECTF	267	44	44	0	0.0	2	79.4	0.00
spectrometer	531	101	101	0	0.0	48	10.4	0.00
sponge	76	44	0	44	3.8	3	92.1	0.65
statlog-australian	690	14	6	8	4.5	2	55.5	0.00
statlog-german	1000	20	7	13	4.2	2	70.0	0.00
statlog-german-numeric	1000	24	24	0	0.0	2	70.0	0.00
statlog-heart	270	13	8	5	2.6	2	55.6	0.00
statlog-vehicle	846	18	18	0	0.0	4	25.8	0.00
tae	151	5	3	2	2.0	3	34.4	0.00
thyroid-disease-new	215	5	5	0	0.0	3	69.8	0.00
tic-tac-toe	958	9	0	9	3.0	2	65.3	0.00
vowel-context	990	10	10	0	0.0	11	09.1	0.00
wine	178	13	13	0	0.0	3	39.9	0.00
zoo	101	16	1	15	2.0	7	40.6	0.00

skewness and kurtosis is difficult and relevant only per each data set and attribute separately so we exclude skewness and kurtosis from the summary comparisons in this section. For numeric attributes we computed p-values of KS tests under the null hypothesis that attribute values from both compared data sets are drawn from the same distribution. We report the percentage of numeric attributes where this hypothesis was rejected at 0.05 level (lower value reported indicates higher similarity). For discrete attributes we compared similarity of value distributions using Hellinger distance. The response variables was excluded from the data sets for this comparison as the similarity of their distributions was enforced by the generator. We report average Hellinger distance over all discrete attribute in each data set in results below.

Clustering: the structure of original and constructed data sets were compared with k-medoids clustering, using ARI (Eq. (2)) as presented in workflow on Fig. 9. The response variables was excluded from the data sets. For some data sets ARI exhibits high variance, so we report the average ARI over 100 repetitions of generating new data.

Classification performance: we compared the predictive similarity of the data sets using classification accuracy of random forests as illustrated in a workflow on Fig. 10. We selected random forests due to the robust performance of this learning algorithm under various conditions [27]. The implementation used comes from R package CORElearn [23]. The default parameters were used: we built 100 random trees with the number of randomly selected attributes in nodes set to square root of the number of attributes. We report 5 x 2 cross-validated performances of models trained and tested on both data sets.

The results of these comparisons are presented in Table 4 for integer encoding of nominal attributes. The column labeled G present number of Gaussian kernels in the constructed generator. Relatively large number of units are needed to adequately represent the training data. Nevertheless the generator construction time (function rbfDataGen) in seconds is low as seen from column labeled t. For measurements we used a single core of Intel i7 CPU running at 2.67Ghz. The time to generate the data (function newdata) was below 1 sec for 1000 instances in all cases, so we do not report it.

The column labeled with $=$ gives the percentage of generated instances exactly equal to the original instances. This mostly happens in data sets with only discrete attributes where the whole problem space is small and identical instances are to be expected. Exception from this are datasets horse-colic, primary-tumor and breast-cancer, where the generators contain majority of Gaussian units with only one activation instance. The reason for this is large number of attributes and consequently a poor generalization of rbfDDA algorithm (note the ratio between the number of instances (and the number of attributes) in Table 3 and the number of Gaussian units in Table 4).

Columns labeled $\overline{\Delta mean}$ and $\overline{\Delta std}$ report average difference in mean and standard deviation for attributes normalized to $[0,1]$. In 38 out of 39 cases the difference is

Table 4: The comparison of original and generated data sets. The columns are: G - the number of Gaussians, t - generator construction time in seconds, $=$ - proportion of generated instances exactly equal to original instances, $\overline{\Delta mean}$ - average difference in means for normalized numeric attributes, $\overline{\Delta std}$ - average difference in standard deviation for normalized numeric attributes, KSp - percentage of p-values below 5% in KS tests comparing matching numeric attributes, $\overline{H}$ - average Hellinger distance for matching discrete attributes, ARI - adjusted Rand index, mXdY - classification accuracy in percents for model trained on data X and tested on data Y (for X, Y: 1-original, 2-generated). The dash means that given comparison is not applicable to the data set.

dataset	G	t	$=$	$\overline{\Delta mean}$	$\overline{\Delta std}$	KSp	$\overline{H}$	ARI	m1d1	m1d2	m2d1	m2d2
annealing	207	8.4	0	-0.009	-0.002	100	0.125	0.287	99	86	94	94
arrhythmia	430	42.0	0	-0.016	-0.020	96	0.059	0.013	71	94	58	100
audiology	138	6.6	73	-	-	-	0.086	0.319	72	80	74	86
automobile	112	2.7	0	-0.018	-0.034	6	0.496	0.256	71	84	75	83
balance-scale	217	6.5	0	-0.023	0.028	100	-	0.293	85	89	84	95
breast-cancer	163	2.9	82	-	-	-	0.404	0.145	72	78	67	85
breast-wdbc	148	5.8	0	-0.024	-0.048	60	-	0.972	96	99	95	99
breast-wisconsin	112	3.7	0	-0.022	0.013	100	-	0.911	96	99	96	99
bridges.version1	69	1.8	0	0.009	-0.027	50	0.456	0.180	62	83	65	96
bridges.version2	72	2.2	1	0.031	0.002	100	0.522	0.432	61	66	42	95
bupa	257	3.7	0	0.012	-0.031	66	-	0.680	72	82	73	88
credit-screening	286	8.4	0	-0.041	-0.040	83	0.290	0.133	87	94	86	95
cylinder-bands	370	11.3	0	-0.003	-0.028	100	0.457	0.444	78	72	74	90
dermatology	105	4.3	0	0.044	-0.015	100	0.399	0.593	97	93	95	95
ecoli	127	2.6	0	-0.026	-0.044	85	-	0.914	84	91	85	92
flags	151	3.5	0	-0.046	-0.078	100	0.297	0.189	61	62	48	94
glass	121	2.3	0	-0.043	-0.023	66	-	0.308	72	84	72	90
haberman	154	2.0	0	-0.054	-0.057	50	0.785	0.235	71	76	74	83
heart-cleveland	216	4.1	0	-0.012	-0.025	66	0.215	0.379	57	73	58	95
heart-hungarian	155	3.9	0	-0.007	-0.031	50	0.320	0.216	83	88	82	94
hepatitis	71	1.3	0	-0.034	-0.050	50	0.041	0.136	83	93	86	97
horse-colic	251	5.0	91	0.049	0.008	42	0.431	0.472	83	97	84	99
house-votes-84	180	4.2	17	-	-	-	0.530	0.162	95	76	56	94
ionosphere	143	3.7	0	0.028	0.023	60	-	0.540	93	99	81	99
iris	24	0.8	0	0.010	-0.018	0	-	0.926	95	94	96	93
labor-negotiations	40	0.6	0	-0.236	-	37	0.343	0.065	90	94	90	95
lymphography	91	2.0	0	-0.011	-0.002	100	0.271	0.238	80	82	81	93
monks-1	234	4.5	97	-	-	-	0.306	0.067	99	72	77	73
monks-2	232	5.7	99	-	-	-	0.310	0.080	79	67	76	70
monks-3	235	5.2	98	-	-	-	0.306	0.086	98	79	60	77
pima-diabetes	481	11.8	0	-0.007	-0.021	87	-	0.626	76	91	79	92
post-operative	61	1.3	0	0.070	-0.014	100	0.360	0.262	67	73	69	96
primary-tumor	280	8.2	92	-	-	-	0.135	0.155	44	50	44	83
promoters	96	2.6	69	-	-	-	0.587	0.202	89	99	50	100
sonar.all	133	5.2	0	-0.018	-0.023	5	-	0.580	77	91	80	93
soybean-large	224	11.6	35	-	-	-	0.237	0.506	92	62	79	95
spect-SPECT	160	4.0	70	-	-	-	0.025	0.699	83	91	85	90
spect-SPECTF	220	5.8	0	0.096	-0.078	97	-	0.240	81	100	79	100
spectrometer	468	33.5	0	-0.012	-0.008	56	-	0.105	49	85	46	99
sponge	21	0.9	7	-	-	-	0.376	0.794	92	96	92	97
statlog-australian	281	8.5	0	-0.042	-0.038	83	0.416	0.708	87	95	87	97
statlog-german	733	24.0	0	-0.026	0.000	100	0.456	0.148	75	86	75	96
statlog-german-n	734	26.4	0	-0.007	0.004	100	-	0.456	76	87	78	91
statlog-heart	137	3.1	0	-0.013	-0.016	62	0.142	0.665	81	91	82	96
statlog-vehicle	552	18.3	0	-0.026	-0.020	100	-	0.917	75	89	74	94
tae	90	36.3	0	-0.012	-0.014	33	0.019	0.489	55	73	64	70
thyroid-new	32	1.2	0	-0.001	-0.022	80	-	0.494	95	96	95	97
tic-tac-toe	845	23.0	89	-	-	-	0.444	0.127	95	66	62	89
vowel-context	302	13.9	0	-0.000	-0.009	0	-	0.387	87	88	90	85
wine	50	1.7	0	-0.018	-0.036	23	-	0.403	96	98	95	96
zoo	24	1.4	0	-0.073	-0.004	100	0.019	0.840	91	96	95	96

below 0.10 and in 33 out of 39 cases it is below 0.05 which shows that moments of the distributions for individual attributes are close to the originals. The distributions of individual attributes are compared with KS test for numeric attributes and with Hellinger distance for discrete attributes. Column labeled *KSp* gives a percentage of p-values below 0.05 in KS tests comparing matching numeric attributes using the null hypothesis that original and generated data are drawn from the same distribution. For most of the data sets and most of the attributes the KS-test detects the differences in distributions. Column labeled $\overline{H}$ presents average Hellinger distance for matching discrete attributes. While for many data sets the distances are low, there are also some data sets where the distances are relatively high, indicating that distribution differences can be considerable for discrete attributes.

The suitability of the generator as a development, simulation, or benchmarking tools in data mining is evidenced by comparing clustering and classification performance. The column labeled ARI presents adjusted Rand index. We can observe that the clustering similarity is considerable for many data sets (high ARI) but there are also some data sets where it is low.

The columns m1d1, m1d2, m2d1, and m2d2 report 5x2 cross-validated classification accuracy of random forest models trained on either original (m1) or generated (m2) data and tested on both original (d1) and generated (d2) data. A general trend observed is that on majority of data sets model trained on original data (m1) performs better on generated data than on the original data (m1d2 is larger than m1d1 for 41 of 51 data sets). This indicates that some of the complexity of the original data is lost in the generated data. This is confirmed also by models built on generated data (m2) which mostly perform better on the generated than on original data (m2d2 is larger than m2d1 in 45 out of 51 data sets). Nevertheless, the generated data can be a satisfactory substitute for data mining in many cases, namely models build on generated data outperforms model built on original data when both are tested on original data in half the cases (m2d1 is larger than m1d1 in 25 cases out of 51, in 26 cases m1d1 is larger and there is 1 draw).

An overall conclusion is therefore that for a considerable number of data sets the proposed generator can generate semi-artificial data which is a reasonable substitute in development of data mining algorithms.

6.1 Binary encoding of attributes

As discussed in Sect. 4.2 we can encode each nominal attribute with a set of binary attributes instead of a single integer attribute and avoid making unjustified assumption about the order of the attribute's values. We report the results of tests (statistical, clustering and classification) using binary encoding in Table 5. As the binary encoding of nominal attributes is used only for nominal non-binary attribute, we report results only for data sets that include at least one such attribute.

The number of Gaussian kernels is mostly larger with binary encoding of attributes (in 19 of 29 cases), and so is generator construction time (in 28 out of 34 cases), but the differences are relatively small (on average 20 more Gaussian units are created, and

Table 5: The comparison of original and generated data sets using binary encoding of nominal data sets. Only data sets containing at least one nominal non-binary attribute are included. The meaning of columns is the same as in Table 4.

dataset	G	t	$=$	$\overline{\Delta mean}$	$\overline{\Delta std}$	KSp	$\overline{H}$	ARI	m1d1	m1d2	m2d1	m2d2
annealing	148	10.3	0	-0.008	0.018	100	0.016	0.848	98	97	98	97
audiology	163	9.7	92	-	-		0.060	0.296	70	84	75	90
automobile	127	4.8	0	-0.023	-0.041	6	0.060	0.391	72	91	74	93
breast-cancer	209	5.7	88	-	-		0.059	0.200	71	89	77	99
bridges.version1	65	2.1	0	0.026	-0.011	50	0.095	0.105	63	83	69	95
bridges.version2	79	2.7	0	0.028	-0.027	0	0.087	0.098	63	87	66	98
credit-screening	338	13.2	0	-0.030	-0.026	83	0.027	0.083	87	94	87	96
cylinder-bands	352	26.5	0	-0.008	-0.031	95	0.113	0.656	78	84	76	90
dermatology	242	14.8	1	0.000	-0.035	0	0.046	0.872	98	98	96	99
flags	164	7.6	0	-0.052	-0.094	100	0.108	0.206	60	87	60	99
haberman	148	3.3	0	-0.024	-0.027	50	0.081	0.762	72	83	76	86
heart-disease-cleveland	195	4.2	0	-0.011	-0.028	66	0.069	0.177	57	75	62	93
heart-disease-hungarian	131	3.3	0	-0.014	-0.027	66	0.137	0.738	83	93	83	93
horse-colic	306	10.1	92	0.051	0.026	85	0.120	0.460	84	94	83	99
house-votes-84	182	8.6	73	-	-		0.052	0.863	96	99	95	99
labor-negotiations	36	1.0	0	-0.246	-	25	0.188	0.056	91	98	89	95
lymphography	111	3.0	0	-0.026	0.006	100	0.067	0.290	81	96	81	96
monks-1	187	4.9	99	-	-		0.011	0.065	97	99	99	98
monks-2	342	8.2	99	-	-		0.013	0.126	81	91	82	89
monks-3	205	5.0	99	-	-		0.014	0.104	98	99	98	99
post-operative	66	1.1	12	-0.042	-0.040	100	0.146	0.493	65	75	69	98
primary-tumor	287	8.7	96	-	-		0.060	0.109	42	70	44	92
promoters	99	8.3	40	-	-		0.409	0.334	85	100	50	100
soybean-large	290	21.4	51	-	-		0.056	0.266	92	95	85	98
sponge	31	3.0	11	-	-		0.169	0.431	92	96	93	97
statlog-australian	342	13.8	0	-0.042	-0.034	100	0.032	0.892	87	96	87	96
statlog-german	841	39.7	0	-0.022	-0.002	100	0.073	0.150	75	89	76	99
statlog-heart	136	3.4	0	-0.030	-0.018	62	0.056	0.730	82	92	84	97
tic-tac-toe	897	32.2	94	-	-		0.095	0.218	95	91	66	100

the generator needs 3.7 seconds more. We compared significance of the differences between integer and binary encodings using Wilcoxon rank sum test at 0.05 level. The binary encoding produces significantly more equal instances, but lower Hellinger distance, higher ARI, and lower difference between m1d1 and m2d1, which all indicate improved similarity to the original data. The differences in numeric attributes were not significant. As a result of these findings we recommend using binary encoding of nominal attributes.

6.2 Correction of estimated spread

In several generators we observed low number of instances activated per kernel, with many kernels being formed around a single instance. For such cases the estimated variance is zero and might cause an overfitting of the training data. We try to alleviate the problem with the parameter defaultSpread in function newdata (see Fig. 5) which, in case of zero estimated variance for certain dimension, replaces this unrealistic value with a (small) constant variance, typically between 0.01 and 0.20. The results for the default value of defaultSpread=0.05 is reported in Table 6.

We compared similarity using the binary encoding of nominal attributes with defaultSpread=0.05 (Table 6) and binary encoding of nominal attributes (Table 5 com-

Table 6: The comparison of original and generated data sets using binary encoding of nominal attributes and value of parameter defaultSpread=0.05. The meaning of columns is the same as in Table 4.

dataset	G	t	=	$\overline{\Delta mean}$	$\overline{\Delta std}$	KSp	$\overline{H}$	ARI	m1d1	m1d2	m2d1	m2d2
annealing	143	12.0	0	-0.098	0.024	100	0.029	0.721	98	92	97	93
arrhythmia	421	40.5	0	-0.019	-0.034	97	0.035	0.218	71	69	56	96
audiology	161	8.3	36	-	-	-	0.052	0.174	71	83	71	83
automobile	125	4.7	0	-0.035	-0.056	53	0.061	0.611	71	67	65	67
balance-scale	212	4.5	0	0.010	0.037	100	-	0.288	84	85	83	88
breast-cancer	208	5.4	66	-	-	-	0.059	0.176	70	86	77	94
breast-cancer-wdbc	151	6.0	0	-0.021	-0.047	97	-	0.951	95	97	95	98
reast-cancer-wisconsin	122	3.9	0	-0.011	0.017	100	-	0.968	96	98	97	98
bridges.version1	68	1.4	0	-0.023	-0.017	25	0.089	0.726	63	74	70	83
bridges.version2	77	1.9	0	-0.039	0.011	0	0.073	0.126	61	82	68	96
bupa	266	4.4	0	-0.035	-0.052	100	-	0.226	69	65	63	73
credit-screening	342	13.4	0	-0.091	-0.063	83	0.032	0.115	87	92	87	94
cylinder-bands	355	27.1	0	-0.015	-0.063	100	0.108	0.602	78	73	73	86
dermatology	257	15.6	0	-0.005	-0.050	0	0.046	0.914	97	97	96	98
ecoli	131	2.8	0	-0.056	-0.041	57	-	0.810	84	85	85	86
flags	167	8.0	0	-0.160	-0.089	100	0.109	0.218	61	74	57	97
glass	123	2.3	0	-0.070	-0.036	67	-	0.383	72	63	69	77
haberman	149	3.1	0	-0.042	-0.030	50	0.055	0.817	71	75	76	80
eart-disease-cleveland	217	4.9	0	-0.031	-0.024	83	0.043	0.537	57	72	61	96
eart-disease-hungarian	124	2.8	0	-0.020	-0.009	67	0.147	0.751	82	91	83	93
hepatitis	83	2.0	0	-0.018	-0.037	0	0.058	0.151	84	91	84	97
horse-colic	314	12.5	84	0.014	-0.016	57	0.123	0.323	84	91	85	93
house-votes-84	180	6.1	53	-	-	-	0.030	0.840	96	99	94	99
ionosphere	160	5.2	0	0.026	0.028	79	-	0.540	93	97	85	98
iris	23	1.5	0	0.005	-0.007	25	-	0.897	95	95	96	98
labor-negotiations	37	0.8	0	-0.207	-	38	0.208	0.048	86	86	88	87
lymphography	117	2.7	0	-0.113	0.008	100	0.086	0.284	80	90	79	93
monks-1	184	4.9	96	-	-	-	0.015	0.060	97	95	97	93
monks-2	345	8.1	94	-	-	-	0.008	0.114	80	88	79	82
monks-3	207	5.5	96	-	-	-	0.011	0.091	98	96	98	95
pima-indians-diabetes	506	12.3	0	-0.021	-0.024	100	-	0.295	76	81	76	84
post-operative	68	1.0	4	0.071	-0.011	100	0.102	0.422	66	76	70	88
primary-tumor	288	8.2	79	-	-	-	0.058	0.109	45	71	44	88
promoters	101	7.5	3	-	-	-	0.335	0.233	88	98	50	100
sonar.all	128	5.2	0	-0.019	-0.023	18	-	0.236	78	91	82	95
soybean-large	300	41.6	28	-	-	-	0.058	0.385	92	94	83	96
spect-SPECT	161	3.8	57	-	-	-	0.020	0.836	82	89	85	89
spect-SPECTF	215	5.7	0	0.111	-0.100	100	-	0.203	81	100	79	100
spectrometer	473	33.1	0	-0.008	-0.032	97	-	0.969	49	57	38	88
sponge	44	3.7	10	-	-	-	0.081	0.966	92	97	93	97
statlog-australian	333	9.6	0	-0.082	-0.053	100	0.030	0.939	87	91	87	93
statlog-german	845	39.1	0	-0.042	0.020	100	0.060	0.253	75	83	76	96
statlog-german-numeric	723	24.5	0	-0.023	0.019	100	-	0.159	75	80	76	83
statlog-heart	139	2.7	0	-0.039	0.001	75	0.056	0.667	81	92	83	94
statlog-vehicle	571	17.5	0	-0.035	-0.031	100	-	0.954	74	72	70	80
tae	90	1.0	0	-0.028	-0.010	33	0.037	0.350	50	54	54	55
thyroid-disease-new	38	1.3	0	-0.021	-0.028	60	-	0.497	96	85	91	90
tic-tac-toe	859	28.6	78	-	-	-	0.090	0.164	95	98	74	99
vowel-context	291	12.4	0	0.004	-0.016	0	-	0.441	88	77	85	76
wine	52	1.8	0	-0.022	-0.035	31	-	0.487	97	96	97	95
zoo	23	1.4	0	-0.030	0.017	100	0.025	0.575	91	94	94	94

bined with Table 4 for data sets missing in Table 5). The Wilcoxon paired rank sum test at 0.05 level shows that this setting produces significantly lower proportion of equal instances, lower difference between means for numeric attributes, and lower the difference between m1d1 and m2d1. Other differences were not significant at this level. The approximation of original data is nevertheless better, so we recommend some experimentation with this setting or using a safe default.

6.3 When RBF-based generator works?

We tried to determine the conditions when RBF based data generation works well. The first hypothesis we tested is whether the success of RBF classification algorithm is related to the success of RBF based data generation. For this we compared the classification performance of rbfDDA with performance of random forests. We selected random forests as it is one of the most successful classifiers, known for its robust performance (see for example [27]). Using 5×2 cross-validation we compared the classification accuracy and AUC of rbFDDA algorithm from RSNNS package [3] with random forest implemented in CORElearn package [23] using the default parameters for both classifiers. Unsurprisingly, random forests produced significantly higher accuracy and AUC. We report results for the accuracy in four left-hand side columns of Table 7. Results for the AUC are highly similar so we skip them.

Unsurprisingly, random forest achieve better accuracy than RBF networks. The difference is significant on 0.05 level for 34 of 51 data sets (for 3 data sets RBF is significantly better, other differences are insignificant).

We tried to identify the main factors affecting the success of proposed RBF based data generator. As a measure of success we use the difference in classification accuracy of models trained on original data (m1) and generated data (m2) and tested on original data (d1). This difference is labeled Δd_1 in Table 7. The factors possibly affecting this indicator are difference in classification accuracies between RBF and RF, number of instances, number of attributes, number of instances per attribute, number of Gaussians in the generator, average number of instances per Gaussian kernel, and number of attributes per Gaussian kernel. These factors are collected in Table 7. In Table 8 we show their correlation with Δd_1.

The Pearson's correlation coefficients indicate largest correlation of performance with difference in classification accuracy between RBF and RF, number of attribute and number of Gaussian kernels. All these factors are indicators of difficulty of the problem for RBF classifier, hinting that the usability of the proposed generator depends on the ability of the learning method to capture the structure of the problem.

We tried to predict the success of the data generator using stepwise linear model with independent variables as above, but it turned out that difference in classification accuracy between RBF and RF is the only variable needed. Other prediction methods were also not successful.

Table 7: Different factors than might influence the quality of RBF-based data generator. The columns are: *RBF* - the classification accuracy of rbfDDA classifier in percents *RF* - classification accuracy of random forests in percents, Δacc - difference in classification accuracy between *RBF* and *RF* in percents, Δd_1 - difference in classification accuracy of m1 on d1 and m2 on d1 in percents, n - number of instances, a - number of attributes, n/a - number of instances per attribute, G - number of Gaussian kernels in the generator, n/G - average number of instances per Gaussian unit, a/G - number of attributes per Gaussian unit.

dataset	*RBF*	*RF*	Δacc	Δd_1	n	a	n/a	G	n/G	a/G
annealing	82	98	16.3	1.69	898	38	23.6	143	6.2	0.26
arrhythmia	61	72	11.5	14.42	452	279	1.6	421	1.0	0.66
audiology	59	72	13.6	-0.92	226	69	3.2	161	1.4	0.42
automobile	59	72	13.4	6.68	205	25	8.2	125	1.6	0.20
balance-scale	89	84	-4.5	0.73	625	4	156.2	212	2.9	0.01
breast-cancer	72	70	-1.8	-7.13	286	9	31.7	208	1.3	0.04
breast-cancer-wdbc	94	95	0.6	0.47	569	30	18.9	151	3.7	0.19
breast-cancer-wisconsin	96	96	-0.2	-0.31	699	9	77.6	122	5.7	0.07
bridges.version1	59	65	6.0	-6.69	106	11	9.6	68	1.5	0.16
bridges.version2	57	63	5.6	-7.07	106	11	9.6	77	1.3	0.14
bupa	62	69	7.0	5.90	345	6	57.5	266	1.3	0.02
credit-screening	83	86	3.0	-0.89	690	15	46.0	342	2.0	0.04
cylinder-bands	69	79	9.8	4.74	540	37	14.5	355	1.5	0.10
dermatology	68	96	28.4	1.28	366	34	10.7	257	1.4	0.13
ecoli	81	84	3.0	-0.38	336	7	48.0	131	2.5	0.05
flags	56	60	3.8	4.43	194	28	6.9	167	1.1	0.16
glass	64	73	8.9	2.89	214	9	23.7	123	1.7	0.07
haberman	68	70	2.2	-4.77	306	3	102.0	149	2.0	0.02
heart-disease-cleveland	56	57	0.7	-4.35	303	13	23.3	217	1.4	0.05
heart-disease-hungarian	81	83	1.9	-0.91	294	13	22.6	124	2.3	0.10
hepatitis	79	83	3.2	-0.52	155	19	8.1	83	1.8	0.22
horse-colic	82	84	2.0	-0.84	368	21	17.5	314	1.1	0.06
house-votes-84	88	95	7.7	1.33	435	16	27.1	180	2.4	0.08
ionosphere	92	93	0.5	7.72	351	34	10.3	160	2.1	0.21
iris	91	93	2.0	-0.33	150	4	37.5	23	6.5	0.17
labor-negotiations	82	88	5.9	-2.42	57	16	3.5	37	1.5	0.43
lymphography	78	83	4.7	0.54	148	18	8.2	117	1.2	0.15
monks-1	76	98	21.3	0.71	556	6	92.6	184	3.0	0.03
monks-2	68	84	16.0	0.31	601	6	100.1	345	1.7	0.01
monks-3	76	98	22.3	0.03	554	6	92.3	207	2.6	0.02
pima-indians-diabetes	74	75	1.8	0.06	768	8	96.0	506	1.5	0.01
post-operative	68	66	-2.8	-4.66	90	8	11.2	68	1.3	0.11
primary-tumor	37	44	6.3	0.67	339	17	19.9	288	1.1	0.05
promoters	64	84	20.0	38.11	106	57	1.8	101	1.0	0.56
sonar.all	71	78	6.4	-3.22	208	60	3.4	128	1.6	0.46
soybean-large	80	93	12.8	9.41	683	35	19.5	300	2.2	0.11
spect-SPECT	80	83	2.4	-3.63	267	22	12.1	161	1.6	0.13
spect-SPECTF	79	80	0.7	1.34	267	44	6.0	215	1.2	0.20
spectrometer	35	49	14.2	10.79	531	101	5.2	473	1.1	0.21
sponge	92	92	-0.2	-1.05	76	44	1.7	44	1.7	1.00
statlog-australian	83	86	3.1	-0.14	690	14	49.2	333	2.0	0.04
statlog-german	70	76	6.2	-1.03	1000	20	50.0	845	1.1	0.02
statlog-german-numeric	70	76	6.5	-1.03	1000	24	41.6	723	1.3	0.03
statlog-heart	80	82	1.6	-1.74	270	13	20.7	139	1.9	0.09
statlog-vehicle	61	74	13.6	4.01	846	18	47.0	571	1.4	0.03
tae	50	53	3.0	-3.63	151	5	30.2	90	1.6	0.05
thyroid-disease-new	94	95	0.4	5.07	215	5	43.0	38	5.6	0.13
tic-tac-toe	73	90	16.8	21.57	958	9	106.4	859	1.1	0.01
vowel-context	84	88	3.8	2.87	990	10	99.0	291	3.4	0.03
wine	94	96	1.9	-0.16	178	13	13.6	52	3.4	0.25
zoo	80	89	9.1	-3.38	101	16	6.3	23	4.3	0.69

Table 8: The correlation coefficient between different factors than might influence the quality of RBF-based data generator and Δd_1 (difference in classification accuracy of m1 on d1 and m2 on d1). The names of the factors are the same as in Table 7.

	RBF	*RF*	Δacc	n	a	n/a	G	n/G	a/G
Δd_1	-0.12	0.12	0.45	0.16	0.38	-0.01	0.27	-0.12	0.21

6.4 Development of big data tools

During development of big data cloud based framework ClowdFlows[13], we wanted to test several classification algorithms which are components of the framework, and also the capabilities and scalability of the framework. Though several examples of public big data problems are freely available, each requires (sometimes tedious) preprocessing and adaptations to the specifics of the problem. As the development already required a significant effort of everyone involved such an additional effort was undesired. The use of rbfDataGen contained in an open-source R package semiArtificial [22] turned out to require little additional work but provided required testing data with desired characteristics. We generated several data sets with different characteristics (varying the number of features, instances, and proportions of classes), which were needed during the development and evaluation of the framework.

7 Conclusions and further work

We present an original and practically useful generator of semi-artificial data which was successfully tested in development of big data tools. The generator captures structure of the problem using RBF classifier and exploits properties of the Gaussian kernel to generate new data similar to the original one. We expect such a tool to be useful in the development and adaptation of data analytics tools to specifics of data sets. Possible other uses are data randomization to ensure privacy, simulations requiring large amounts of data, testing of big data tools, benchmarking, and scenarios with huge amounts of data.

We developed a series of evaluation tools which can provide an estimate of generator's performance for specific data sets. Using a large collection of UCI data sets we were able to show that the generator was in most cases successful in generating artificial data similar to the original. The success of the generator is related to the success of RBF classifier: where the RBF can successfully capture the properties of the original data, the generator based on RBF will also be successful, and vice versa. Nevertheless we were unable to create a successful prediction model for the quality of the generator. The user is therefore advised to use the provided evaluation tools on the specific data set. The provided results shall provide a good indication on the usability of the generated data for the intended use. The proposed generator together with statistical, clustering and classification performance indicators was turned into an open-source R package semiArtificial [22].

In the future we plan to extend the generator with new modules using different learning algorithms to capture data structure and generate new data. An interesting approach would also be a rejection approach which uses probability density estimates based on various learning algorithms.

Acknowledgments

Many thanks to Petr Savicky for interesting discussions on the subject and for the preparation of UCI data sets used in this paper. The author was supported by the Slovenian Research Agency (ARRS) through research programme P2-0209.

References

[1] K. Bache and M. Lichman. UCI machine learning repository, 2014. URL `http://archive.ics.uci.edu/ml`.

[2] H. D. Bandara and A. P. Jayasumana. On characteristics and modeling of p2p resources with correlated static and dynamic attributes. In *IEEE Global Telecommunications Conference (GLOBECOM 2011)*, pages 1–6. IEEE, 2011.

[3] C. Bergmeir and J. M. Benítez. Neural networks in R using the Stuttgart neural network simulator: RSNNS. *Journal of Statistical Software*, 46(7):1–26, 2012.

[4] M. R. Berthold and J. Diamond. Boosting the performance of RBF networks with dynamic decay adjustment. In *Advances in Neural Information Processing Systems*, pages 521–528. MIT Press, 1995.

[5] P. A. Ferrari and A. Barbiero. Simulating ordinal data. *Multivariate Behavioral Research*, 47(4):566–589, 2012.

[6] C. Fraley, A. E. Raftery, T. B. Murphy, and L. Scrucca. mclust version 4 for R: Normal mixture modeling for model-based clustering, classification, and density estimation. Technical Report No. 597, University of Washington, Department of Statistics, 2012. URL `http://CRAN.R-project.org/package=mclust`. R package, version 4.2.

[7] J. C. Gower. A general coefficient of similarity and some of its properties. *Biometrics*, pages 857–871, 1971.

[8] H.-G. Han and J.-F. Qiao. Adaptive computation algorithm for RBF neural network. *IEEE Transactions on Neural Networks and Learning Systems*, 23(2): 342–347, 2012.

[9] W. Härdle and M. Müller. Multivariate and semiparametric kernel regression. In M. G. Schimek, editor, *Smoothing and regression: approaches, computation, and application*, pages 357–392. John Wiley & Sons, 2000.

[10] C. Hennig. *fpc: Flexible procedures for clustering*, 2013. URL http://CRAN.R-project.org/package=fpc. R package, version 2.1-7.

[11] L. Hubert and P. Arabie. Comparing partitions. *Journal of Classification*, 2(1): 193–218, 1985.

[12] L. Kaufman and P. J. Rousseeuw. *Finding groups in data: an introduction to cluster analysis*. John Wiley & Sons, 1990.

[13] J. Kranjc, R. Orač, V. Podpečan, M. Robnik-Šikonja, and N. Lavrač. ClowdFlows: Workows for big data on the cloud. Technical report, Jožef Stefan Institute, Ljubljana, Slovenia, 2014.

[14] M. Maechler, P. Rousseeuw, A. Struyf, M. Hubert, and K. Hornik. *cluster: Cluster Analysis Basics and Extensions*, 2013. URL http://CRAN.R-project.org/package=cluster. R package, version 1.14.4.

[15] P. Mair, A. Satorra, and P. M. Bentler. Generating nonnormal multivariate data using copulas: Applications to SEM. *Multivariate Behavioral Research*, 47(4): 547–565, 2012.

[16] J. Moody and C. J. Darken. Fast learning in networks of locally-tuned processing units. *Neural Computation*, 1(2):281–294, 1989.

[17] R. B. Nelsen. *An introduction to copulas*. Springer, 1999.

[18] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2013. URL http://www.R-project.org/.

[19] W. M. Rand. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association*, 66(336):846–850, 1971.

[20] D. L. Reilly, L. N. Cooper, and C. Elbaum. A neural model for category learning. *Biological cybernetics*, 45(1):35–41, 1982.

[21] B. D. Ripley. *Stochastic Simulation*. Wiley, New York, 1987.

[22] M. Robnik-Sikonja. *semiArtificial: Generator of semi-artificial data*, 2014. URL http://CRAN.R-project.org/package=semiArtificial. R package version 1.2.0.

[23] M. Robnik-Šikonja and P. Savicky. *CORElearn - classification, regression, feature evaluation and ordinal evaluation*, 2012. URL http://CRAN.R-project.org/package=CORElearn. R package version 0.9.39.

[24] J. Ruscio and W. Kaczetow. Simulating multivariate nonnormal data using an iterative algorithm. *Multivariate Behavioral Research*, 43(3):355–381, 2008.

[25] P. Savicky. *readMLData: Reading machine learning benchmark data sets from different sources in their original format*, 2012. URL `http://CRAN.R-project.org/package=readMLData`. R package, version 0.9-6.

[26] W. N. Venables and B. D. Ripley. *Modern Applied Statistics with S*. Springer, New York, 4th edition, 2002.

[27] A. Verikas, A. Gelzinis, and M. Bacauskiene. Mining data with random forests: A survey and results of new tests. *Pattern Recognition*, 44(2):330–349, 2011.

[28] N. X. Vinh, J. Epps, and J. Bailey. Information theoretic measures for clusterings comparison: is a correction for chance necessary? In *Proceedings of the 26th International Conference on Machine Learning, ICML 2009*, pages 1073–1080, 2009.

[29] T. Xie, H. Yu, J. Hewlett, P. Rozycki, and B. Wilamowski. Fast and efficient second-order method for training radial basis function networks. *IEEE Transactions on Neural Networks and Learning Systems*, 23(4):609–619, 2012.

[30] A. Zell, G. Mamier, M. Vogt, N. Mache, R. Huebner, S. Doering, K.-U. Herrmann, T. Soyez, M. Schmalzl, T. Sommer, A. Hatzigeorgiou, D. Posselt, T. Schreiner, B. Kett, G. Clemente, J. Wieland, and J. Gatter. SNNS: Stuttgart neural network simulator. User manual, version 4.2. Technical report, University of Stuttgart and University of Tuebingen, 1995.